

# AN ONBOARD DATA ANALYSIS METHOD TO TRACK THE SEASONAL POLAR CAPS ON MARS

Kiri L. Wagstaff<sup>1</sup>, Rebecca Castaño<sup>1</sup>, Steve Chien<sup>1</sup>, Anton B. Ivanov<sup>1</sup>, Erik Pounders<sup>1</sup>, and Timothy N. Titus<sup>2</sup>

<sup>1</sup>Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA,  
Email: firstname.lastname@jpl.nasa.gov

<sup>2</sup>United States Geological Survey, 2255 N. Gemini Dr., Flagstaff, AZ 86001, USA, Email: ttitus@usgs.gov

## ABSTRACT

The Martian seasonal CO<sub>2</sub> ice caps advance and retreat each year. They are currently studied using instruments such as the THERMAL EMISSION Imaging System (THEMIS), a visible and infra-red camera on the Mars Odyssey spacecraft [1]. However, each image must be downlinked to Earth prior to analysis. In contrast, we have developed the Bimodal Image Temperature (BIT) histogram analysis method for onboard detection of the cap edge, before transmission. In downlink-limited scenarios when the entire image cannot be transmitted, the location of the cap edge can still be identified and sent to Earth. In this paper, we evaluate our method on uncalibrated THEMIS data and find 1) agreement with manual cap edge identifications to within 28.2 km, and 2) high accuracy even with a reduced context window, yielding large reductions in memory requirements.

## 1. INTRODUCTION

Like the Earth, Mars experiences significant seasonal weather patterns. One result of these changes is the presence of CO<sub>2</sub> ice caps at both poles that advance and recede seasonally. The seasonal cycling of CO<sub>2</sub> from the atmosphere to the polar caps (condensation) and back to the atmosphere (sublimation) significantly alters the distribution of mass on the planet, to the extent that it is possible, even from Earth, to observe the oscillation of the center of gravity of Mars [2]. We are interested in automatically tracking the motion of the polar caps over time so that we can better understand the processes at the north and south poles as well as any interannual changes in polar cap behavior.

Since the polar cap stands out as distinctly colder than the rest of Mars, an ideal way to track it is to use an infra-red camera in Mars orbit. In this study, we have used data collected by the Thermal Emission Imaging System (THEMIS) instrument [1] onboard the Mars Odyssey spacecraft. Mars Odyssey started its science orbits in

February of 2002 and has made observations that cover two northern Martian winters to date.

In this paper, we present the Bimodal Image Temperature (BIT) histogram analysis method and its application to THEMIS images of the north polar region of Mars. We compare its detections to two standards: 1) a model of the polar cap generated by an independent analysis of contemporary data collected by the Thermal Emission Spectrometer (TES), and 2) manual annotations of the same THEMIS data produced by visual inspection of the temperature profile. We find very good agreement with both standards. In Section 6.2, we present the results of using the BIT method in a simulated onboard scenario, where THEMIS continually collects data and analyzes a cached sliding window of data. We find that recall is affected to a larger degree than precision when window size is reduced. This effect may be tolerable if a reduction in the computational and memory requirements is needed.

We conclude in Section 7 with a discussion of the advantages of using BIT over a TES-based model or manual annotations, and we describe how BIT could be used onboard THEMIS to provide continuous monitoring of the Martian polar caps.

## 2. BACKGROUND ON THEMIS DATA

Before discussing our analysis method, we provide an overview of the data collected by THEMIS. The THEMIS instrument contains two subsystems: Infrared (IR) and Visible (VIS). The IR camera has 10 filters, from 6 to 15  $\mu\text{m}$ , and it images the surface with 100-m/pixel resolution.

Previous analysis of the Martian polar caps has relied on TES data, which has a per-pixel spatial resolution of 3 km. THEMIS provides a 30-fold increase in spatial resolution over TES data and thereby promises to yield much more precise identification of the edge of the polar cap. We focus on data collected by THEMIS band 9 (12.57  $\mu\text{m}$ ), which most clearly shows temperature dif-

ferences in the image. Each THEMIS IR image is 320 pixels (32 km) wide and a variable number (3600 to 14, 352) of pixels long. To study the seasonal north polar cap recession, we identified images of the northern latitudes that ranged from  $L_s$  350° (early winter) to 70° (late spring). Since our goal is to demonstrate the feasibility of on-board analysis, we worked solely with the raw, uncalibrated (EDR) data.

### 3. THE BIT HISTOGRAM ANALYSIS METHOD

Images that contain the polar cap edge have an associated temperature histogram that is bimodal, since both “cold” (polar cap) and “warm” (not polar cap) pixels are present. When this is the case, we dynamically identify the cutoff between cap and non-cap pixels, then locate the image line that best matches the transition between the two. We next describe BIT in more detail.

#### 3.1. The BIT Algorithm

*Step 1: Calibration [Optional].* Because we are working with uncalibrated data, we do not know the “true” temperature of each pixel. However, our method does not rely on absolute temperatures and, in fact, can be applied without any calibration. However, we gain a slight improvement in precision by performing a fast approximate calibration to help select the most appropriate cutoff. We pseudo-calibrate each pixel  $i$  in the image by converting the raw digital number,  $DN_i$ , to a temperature,  $T_i$ , as suggested by Josh Bandfield of ASU:

$$x = (DN_i - o \times g) \times g / 16 \quad (1)$$

$$T_i = 101.85 \times \log_{10}(x) - 223.3. \quad (2)$$

where  $o$  and  $g$  are the instrument offset and gain parameters, provided in the header of the data file.

*Step 2: Temperature Histogram.* We construct a histogram of all of the calibrated temperature values in the image. Each histogram bin is 2 K wide, and the histogram ranges from 130 to 270 K. For onboard use, we will instead pre-convert the histogram bin temperatures into DNs, and no onboard calibration will be necessary. Fig. 1 shows the histogram for image I09640013.

*Step 3: Dynamic Thresholding.* We identify the characteristic “dip” (local minimum) between the two temperature modes, and select the corresponding temperature,  $T'$ , as the threshold that distinguishes the polar cap from non-cap pixels. More specifically, we first identify the left and right peaks as local maxima in the histogram. We then identify the minimal point between them as the appropriate temperature threshold. In Fig. 1,  $T' = 172$  K. Due to the season, some images include regions that the sun does not illuminate, resulting in a third (extremely cold) mode. In such a case, we select the “dip” closest to 170 K. Finally, we filter the detections by requiring that  $T'$  be in the range [160, 210] K.




Figure 1. Temperature histogram for pixels in THEMIS image I09640013 (longitude 268.42E,  $L_s = 350.61^\circ$ ).

*Step 4: Cap Edge Identification.* The CO<sub>2</sub> cap is not a discrete phenomenon, with an abrupt “edge”. Instead, it grows thinner from north to south, eventually becoming a thin layer of CO<sub>2</sub> frost, then isolated frost deposits, and then disappearing completely. Therefore, we define the cap “edge” as the point at which only 50% of the surface is covered in frozen CO<sub>2</sub>. BIT applies the temperature threshold to the original image by marking each pixel that is colder than  $T'$  as belonging to the polar cap and each pixel that is warmer than  $T'$  as “non-cap”. We then proceed from north to south and examine each line of the image, halting when we find a line that is less than 50% “cap”. For image I09640013, BIT finds the cap edge at latitude 59.80 degrees north. This is shown visually as the red line in Fig. 2.

#### 3.2. Complexity Analysis

The BIT algorithm strikes a compromise between simplicity (for potential use onboard THEMIS) and adaptability (to be able to run on uncalibrated data). Overall, the algorithm is linear in the number of pixels in the image,  $n$ . More specifically, the cost of each step is as follows<sup>1</sup>, plus constants:

Step	Action	Operations
1	Calibrate	4 per pixel
2	Update bin count	1 per pixel
3	Find histogram “dip”	1 per bin
4a	Mark as “cap” or “non-cap”	1 per pixel
4b	Sum “cap” pixels in line	1 per pixel
4c	Check if < 50% “cap”	1 per line
Total	7 per pixel, 1 per bin, 1 per line	

<sup>1</sup>Note that we can avoid the  $\log_{10}()$  calculation in the calibration step by pre-converting each histogram bin from temperature  $T$  to an exponential scale.




Figure 2. Uncalibrated temperature data (band 9) from THEMIS image I09640013 (contrast-enhanced), cropped to show only  $58.36^{\circ}\text{N}$  to  $62.33^{\circ}\text{N}$ . The cap edges identified by the TES model (blue, with cyan for  $\pm 1.4$  degrees error estimate), BIT method (red), and manual annotation (green) are shown. The image width is 32 km.

For example, image I09640013 in Fig. 1 is 320 pixels wide by 14,352 lines long, and the histogram we use contains 70 bins. The total number of operations is roughly 32,000,000, and the entire image must be stored in memory (4.4 MB). As we will show later, moving to a sliding window operational scenario can greatly reduce the computation and memory requirements for each analysis performed.

#### 4. RELATED WORK: TES MODEL

The Thermal Emission Spectrometer (TES), onboard Mars Global Surveyor, is also an infra-red camera in Mars orbit. TES observes at wavelengths ranging from 6 to  $50\text{ }\mu\text{m}$ . Although TES has much lower spatial resolution than THEMIS, its temperature observations are much more reliably calibrated. Therefore, we have evaluated BIT's detections against a model derived from contemporary TES observations. This model is a 51-coefficient Fourier fit to cap edge locations identified in 60-km binned TES data, with a 1-sigma error estimate of 1.4 degrees of latitude [3]. For image I09640013, the TES model predicts that the cap edge is at 60.87 degrees north (see Fig. 2, blue and cyan lines), which is 1.07 degrees farther north than BIT's detection and within the margin of error for the TES model.

#### 5. MANUAL ANNOTATIONS

As a second source of independent validation, a member of the team who was not involved in the algorithm development manually annotated 435 THEMIS images. Rather than looking at the temperature histogram, these annotations were generated after examining each image's temperature profile, which includes the mean temperature value for each line in the image (i.e., cross-track averaging). In particular, we examined the temperature profile of each image to identify the beginning and end of the defrosting zone.

For example, the temperature profile for image I09779015 is shown in Fig. 3. This profile shows the characteristic shape we observe in THEMIS images that contain the edge of the  $\text{CO}_2$  cap: a sigmoid curve that transitions from a low temperature compatible with the presence of  $\text{CO}_2$  ice to a temperature that is too warm to support  $\text{CO}_2$  ice. The beginning of the defrosting zone, as annotated, occurs at about line 6200, at a temperature of 160 K, and it ends near line 9600, at 185 K. Our manual annotator identified the beginning and end of this zone to the nearest 100 lines. Therefore, each annotation is specified  $\pm 100$  km (1 line = 100 m). We interpret the midpoint, at line 7900, as a first approximation to the edge of the cap. The same method was used to derive the green line in Fig. 2, which is where the manual annotations predict the cap edge to be.




Figure 3. Temperature profile for THEMIS image I09779015 (longitude 220.57E,  $L_s = 356.44^\circ$ ), with cross-track averaging. The north end of the image is to the left, and the south end is to the right. Vertical lines indicate the manual annotations of the beginning and end of the defrosting zone.

A natural question to ask is to what degree the manual annotations and the TES model predictions agree. For each of the 435 images, we compared the line indicated by the midpoint between the annotated beginning and end of the defrosting zone against the TES model’s predictions for that longitude and  $L_s$ . We found high agreement in terms of deciding which images contained the polar cap (426 of 435), but less agreement about the exact location of the cap. The mean deviation between the manual annotations and the TES model was 2.07 degrees of latitude, or about 124 km, with a strong southward bias; that is, the manual annotations tended to indicate that the polar cap edge was further south than what the TES model would predict. In the rest of this paper, we will evaluate against both standards (TES model and manual annotations), but we will rely more heavily on our comparison with the manual annotations, as they were derived from the same THEMIS data that we provide to BIT.

## 6. EXPERIMENTAL RESULTS

To evaluate the performance of the BIT method, we used it to analyze the 435 THEMIS images in their raw, uncalibrated (EDR) format. We evaluated BIT in two operational scenarios: image-based and continuous (windowed) mode. In each scenario, we used two evaluation metrics. First, we assessed the degree of accuracy in terms of identifying which images contain the CO<sub>2</sub> cap edge, using recall ( $R$ ) and precision ( $P$ ) measures. Let  $I_c$  be the set of images that contain the cap edge (according to the TES model or the manual annotations) and  $BIT_c$  be the set of images BIT labels as containing the cap edge. Then

$$R = \frac{|I_c \cap BIT_c|}{|I_c|} \quad (3)$$

and

$$P = \frac{|I_c \cap BIT_c|}{|BIT_c|}. \quad (4)$$

Recall is the percentage of  $I_c$  that BIT found, and precision is the percentage of  $BIT_c$  that were accurate. False positives are images that lack the CO<sub>2</sub> cap edge but that BIT labels as having the edge present, while false negatives are images that BIT incorrectly labels as not having the cap edge. Note that, in a scenario where BIT’s decisions aid in prioritizing data for downlink, we would rather have false positives than false negatives, so that we can avoid missing any images that contain the cap edge (at the cost of possibly downlinking images that lack it).

The second evaluation metric we used was an assessment of the accuracy of the detected cap edge, when one was detected. We calculated the deviation between BIT’s cap edge detections and those provided by the TES model and the manual annotations, in terms of kilometers and/or degrees of latitude. One degree of latitude on Mars is approximately 60 km.

### 6.1. Results for Image-Based Analysis

THEMIS currently operates in a targeted, image-based fashion. In this operational mode, we assess the accuracy of BIT’s identification of which images contain the polar cap. For those images, we further assess the accuracy with which BIT identifies the location of the cap edge, compared to both the TES model’s predictions and our manual annotations of the images.

In terms of identifying which images contain the CO<sub>2</sub> polar cap edge, we find good agreement with both the TES model (96.3%) and with the manual THEMIS annotations (93.5%); see Table 1. BIT’s precision as measured against both standards is very high, reflecting the small number of false positives that BIT detects. Recall is somewhat lower in both cases, due to the larger number of false negatives. As noted above, in a downlink prioritization scenario, we would prefer to have more false positives than false negatives. We are currently looking into ways to bias the algorithm in this direction.

Table 1. Analysis of the agreement between BIT and two standards, the TES model and manual THEMIS annotations, in identifying which images contain the CO<sub>2</sub> polar cap edge (of 435 total images).

Standard	Precision	Recall	False pos	False neg	Agreement
TES model	97.2%	92.0%	4	12	419 (96.3%)
Manual annotations	94.3%	86.4%	8	21	406 (93.5%)




Figure 4. Comparison of BIT CO<sub>2</sub> ice cap detections (open circles) against the TES model (closed circles) and manual annotations (plus signs), for  $L_s = 340^\circ$  to  $360^\circ$  (late winter / early spring). Detections for the same image are connected by vertical lines. Dashed lines show the TES model’s range of predictions, including the margin of error.

For 80 images where the cap edge is detected, we also evaluated the degree to which the location BIT identified as the cap edge matched both standards. Fig. 4 shows all BIT detections between  $L_s$   $340^\circ$  to  $360^\circ$ , which corresponds to late winter / early spring in the northern hemisphere. The three data points (BIT detection, manual annotation, and TES model prediction) associated with each THEMIS image are connected by a vertical line. The dashed lines indicate the range of TES model predictions for that time period, including the 1-sigma error estimate, as a reference. We see that BIT’s detections and the manual annotations tend to fall within this range. Note that there is one significant outlier, where the BIT method detected the cap edge to be at a latitude in excess of 85 degrees. This detection could be easily discarded due to the time of year at which the image was collected.

Overall, the mean deviation between BIT detections and the TES model was 1.21 degrees (about 72 km, and within the TES model margin of error), while the mean deviation between BIT and the manual annotations was just 28.2 km. Fig. 5 shows histograms of the deviations we observe between BIT detections and each standard. Note that Fig. 5(a) contains vertical lines indicating the 1-sigma error estimate, or the margin of error. We find that the ma-

jority (67.9%) of BIT’s detections fall within this region, with the exceptions tending to occur later in the spring as the cap edge recedes further north. We also observe a bias in both cases; BIT tends to detect the cap edge slightly further south than where the TES model predicts it to be, and slightly north of where the manual annotations indicate that it is. This is consistent with our comparison of the TES model against the manual THEMIS annotations. We also find that the manual annotations histogram, in Fig. 5(b), is much “tighter” than that for the TES model, in Fig. 5(a), indicating a generally higher level of agreement. In this case, 92.6% of BIT’s detections fall within 1.4 degrees of the manual annotations. This observation reflects the fact that BIT’s detections and the manual annotations were generated from the same data source and confirms that BIT is detecting a valid phenomenon.

**Error analysis.** We carefully surveyed the mistakes BIT made to obtain a good understanding of its strengths and weaknesses. The simple, heuristic nature of the algorithm provides efficient operation, but it can sometimes result in very wrong decisions. For example, the biggest class of errors arose from images such as I10631008. The temperature profile for this image is shown in Fig. 6, where we see that only the beginning of the defrosting zone is present (line 8200); the image does not fully capture the end of the defrosting zone. Because of the incomplete imaging of the defrosting zone, the temperature histogram (Fig. 7) is biased to be slightly colder than we would otherwise expect. In this case, BIT identified the temperature threshold has  $T' = 172$  K, a very reasonable and common value. The temperature profile also displays an intriguing temporary rise in temperature around lines 3500 to 5000, where several pixels exhibit temperatures above 172 K. The result is that BIT detects the cap edge prematurely, and much too far north, due to these unusually warm pixels. Images such as this one are the outliers in Fig. 5(b), with deviations to the far right on the histogram. If BIT also examined the temperature profile, it should be possible to avoid this kind of error.

## 6.2. Results for Windowed Analysis

THEMIS will soon be moving to a continuous monitoring mode, where it will collect data continuously. We simulated this mode by running a sliding window, of a specified length, down each image and providing BIT only the lines that fall into the window. We experimented with several different values for this context window size,  $w$ .

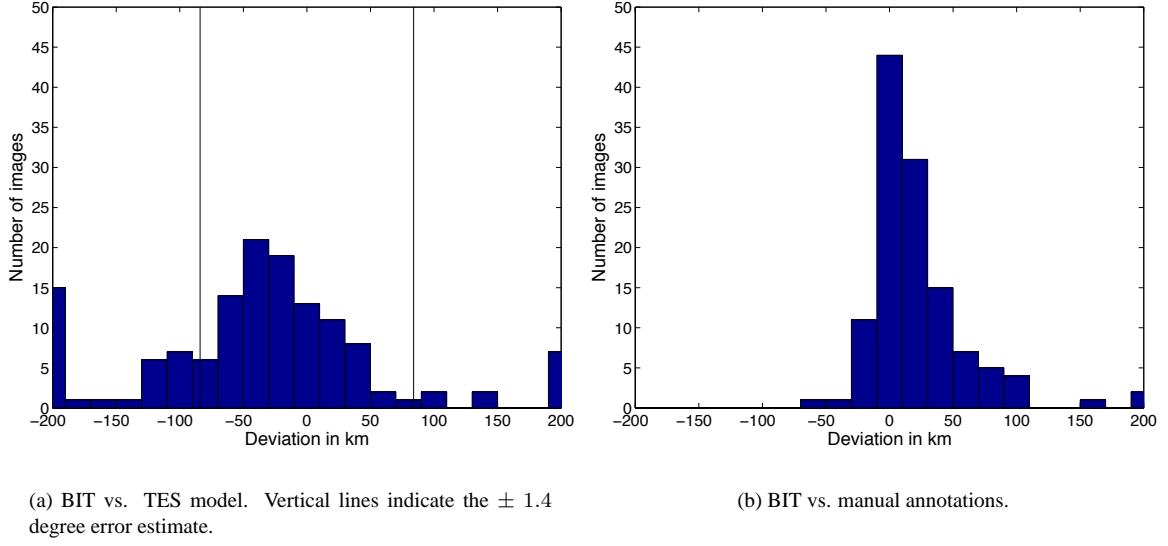


Figure 5. Histogram of deviations (in km) between BIT detections of the polar cap edge and both the TES model and manual THEMIS annotations. Positive deviations indicate that BIT detected the cap farther north.




Figure 6. Temperature profile for THEMIS image I10631008 (longitude 131.08E,  $L_s = 30.09^\circ$ ), with cross-track averaging. The north end of the image is to the left, and the south end is to the right.






Figure 7. Temperature histogram for pixels in THEMIS image I10631008 (longitude 131.08E,  $L_s = 30.09^\circ$ ).

We expect to see a reduction in both recall and precision as the window size is reduced. However, a smaller window requires less onboard memory. In this section, we analyze this tradeoff.

Our simulation was conducted by converting each image into a set of smaller sub-images of size  $w$ . The first sub-image was composed of image lines 1 through  $w$ . Advancing  $w/2$  lines, the next sub-image was composed of lines  $w/2 + 1$  through  $w + w/2$ . This process continued until we reached the end of the original image (the final sub-image generated may be less than  $w$  lines long). Smaller window sizes resulted in a larger number of sub-images that were generated.

We found that performance exhibited a strong dependence on the context window size. Fig. 8 shows the precision and recall we observed when comparing BIT to both the TES modal and the manual annotations for window sizes ranging from 256 to 4096 lines. We find that precision is affected less than recall in terms of the impact of a smaller window. That is, a smaller window size makes it more likely that BIT will miss the cap edge, due to a paucity of data, but tends not to increase the number of false detections as much. We can also see that BIT continues to agree more strongly with the manual annotations (Fig. 8(b)) than with the TES model (Fig. 8(a)).

It is also important to consider the savings in memory requirements provided by the reduced context window size. Fig. 9 shows the amount of memory (in KB) required to process a single context window. Although a smaller window size results in more individual windows to process, each one can be accomplished faster and within tighter memory constraints. For example, the fixed memory size required to search for the polar cap in THEMIS images, using a window size of 2048 lines, is 625 KB. This is less than 20% of the memory required to process an entire image in one pass. For comparison, it would take 3.3 MB to process image I09779015, which is 10,768 lines long. As mentioned earlier, image I09640013, which is 14,352




Figure 9. The memory required for various context window sizes as specified to BIT for continuous monitoring of the Martian surface.

lines long, requires 4.4 MB. Given onboard computing constraints, it is possible to adjust BIT's configuration, via the context window size, to make it feasible for onboard use (at some cost to accuracy).

## 7. CONCLUSIONS

In this paper, we have presented the BIT (Bimodal Image Temperature) histogram analysis method, which can automatically identify the edge of the  $\text{CO}_2$  polar cap in THEMIS images of Mars. In evaluation over 435 images, we have shown that this method has very good agreement both with a model derived from contemporaneous observations by a different instrument (TES on Mars Global Surveyor) and with manual annotations of the THEMIS image by an independent party. The agreement is highest with the manual annotations (within 28.2 km), which increases our confidence in the results. We have identified some cases where BIT makes incorrect decisions and are in the process of developing improvements to address those cases.

Given that THEMIS will be moving to a continuous monitoring mode, we also evaluated BIT in a windowed scenario, where it has access only to data from a fixed number of image lines, rather than the full image. We find that recall and precision are affected by  $w$ , the window size. Our assessment of the algorithm's sensitivity can be used to identify an appropriate tradeoff point between accuracy and memory consumption.

BIT is now ready for onboard use, and it is both accurate and efficient enough for that environment. It was specifically designed to operate with uncalibrated data in the form that it is collected onboard THEMIS. We are currently in the process of converting this algorithm into low-level code suitable for validation in the Mars Odyssey testbed.




Figure 8. Precision and recall for BIT in continuous monitoring mode, as a function of window size,  $w$ .

## ACKNOWLEDGMENTS

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. It was funded by the New Millenium Program and the Mars Odyssey Participating Scientist program. We would like to thank the THEMIS team for their assistance, and Nghia Tang for his contributions in the development and validation of the BIT method.

## REFERENCES

1. Christensen, P. R., Bandfield, J. L., Bell, J. F., Gorelick, N., Hamilton, V. E., Ivanov, A., Jakosky, B. M., Kieffer, H. H., Lane, M. D., Malin, M. C., McConnochie, T., McEwen, A. S., McSween, H. Y., Mehall, G. L., Moersch, J. E., Neelson, K. H., Rice, J. W., Richardson, M. I., Ruff, S. W., Smith, M. D., Titus, T. N., and Wyatt, M. B. Morphology and composition of the surface of Mars: Mars Odyssey THEMIS results. *Science*, 300:2056–2061, June 2003.
2. Smith, D. E. and Zuber, M. T. Seasonal changes in the icecaps of mars from laser altimetry and gravity. In *Proceedings of the 13th International Workshop on Laser Ranging: Science Session*, 2002.
3. Titus, T. N. Mars polar cap edges tracked over 3 full Mars years. In *Proceedings of the 36th Lunar and Planetary Science Conference*, March 2005.